arXiv:1608.03535v2 [cs.PL] 22 Sep 2016

Intuitionistic Logic Programming for SQL (Extended Abstract)

Fernando Sáenz-Pérez

Dept. Ingeniería del Software e Inteligencia Artificial
Declarative Programming Group
Universidad Complutense de Madrid, Spain
fernan@sip.ucm.es

Abstract. Intuitionistic logic programming provides the notion of embedded implication in rule bodies, which can be used to reason about a current database modified by the antecedent. This can be applied to a system that translates SQL to Datalog to solve SQL `WITH` queries, for which relations are locally defined and can therefore be understood as added to the current database. In addition, assumptions in SQL queries as either adding or removing data can be modelled in this way as well, which is an interesting feature for decision-support scenarios. This work suggests a way to apply intuitionistic logic programming to SQL, and provides a pointer to a working system implementing this idea.

Keywords: Intuitionistic Logic Programming, SQL, Datalog

1 Introduction

SQL is the *de facto* relational database query language that stands still [2] despite the advent of new trends as Big Data, NoSQL, RDF stores and others. It builds upon the Codd's [8,9] seminal relational data model accompanied by an algebra and calculus to operate on data. Former proposals such as [12] better provide a formal framework for current SQL implementations. As a query language, SQL can be well understood from Codd's tuple relational calculus but also from logic programming (in particular, [23] includes equivalences between relational operations and logic rules). However, among other features beyond the original relational model, SQL provides the notion of temporary view defined in `WITH` clauses (as described in Section 2), whose definition is available only to the query in which it occurs [21]. This is no longer representable either in relational formal languages or directly in logic programming.[1]

Here is when intuitionistic logic programming may come at help to providing first-class citizen semantics: Approaches as [10,14,15,4,13] fit into this logic, an extension of logic programming including in particular embedded implications. Adding negation to intuitionistic logic programming might develop paradoxes

[1] Obviously, logic programming implementations provide general-purpose languages as Prolog that can emulate a temporary definition.

which are circumvented in [5] by dealing with two kind of implications: for rules ($\leftarrow$) and for goals ($\Leftarrow$, i.e., an embedded implication). Whereas in the formula $A \leftarrow B$, the atom B is "executed" for proving A, in the formula $A \Leftarrow B$, the atom B is "assumed" to be true for proving A. Hypothetical Datalog [3,4,5,6] incorporated this logic and has been a proposal thoroughly studied from semantic and complexity point-of-views. The work [17] (recalled in Section 3) presented an extended (w.r.t. [3,4,5,6]) intuitionistic setting along with an implementation in the deductive system DES [20], in which a rule is accepted in the antecedent of an embedded implication, and not only facts as in [5].

Driven from the need for supporting a broader subset of SQL in this system, we show how to take advantage of the intuitionistic embedded implication to model `WITH` SQL queries in a logic setting, an application which has not been proposed to the best of our knowledge so far. Thus, as shown in Section 4, it is possible to have such SQL queries translated into Datalog, and can be therefore processed by a deductive engine. But Hypothetical Datalog is powerful enough to even apply the same technique to model assumptions in SQL queries, with the (non-standard) clause `ASSUME`. This clause enables both positive and negative assumptions on data, as shown in Section 5, which are useful for modelling "what-if" scenarios. Finally, we present the deductive system DES at work with examples of `WITH` and `ASSUME` queries in Section 6. Our approach is also useful for connecting with external relational database systems which cannot process these clauses. DES then behaves as a front-end capable of processing either novel or unsupported features in such systems.

2 The SQL `WITH` Clause

Typically, complex SQL queries are broken-down for applying the *divide-and-conquer* principle as well as for enhancing readability and maintenance. Introducing intermediate views with `CREATE VIEW` statements is in the order of the day, but this might neither be recommendable (making these views observable for other users) nor possible (only certain users with administration permissions are allowed to create views). The `WITH` clause provides a form of encapsulation in SQL by locally defining those broken-down views, making their realms to pertain to the context of a given query. Next, the syntax of a query Q including this clause is recalled:

```
WITH R1 AS SQL1,
     ...,
     Rn AS SQLn
SQL
```

where each R_i is a temporary view name defined by the SQL statement SQL_i, and which can be referenced only in SQL, the ultimate query that builds the outcome of the query Q. This query can be understood as a relation with name R and defined with the DDL statement `CREATE VIEW` R `AS` Q.

With respect to the semantics of an SQL query, we recall and adapt the notation in [7] which in turn is based on [12]. A *table instance* is a multiset of

facts (following logic programming instead of relational databases). A *database instance* Δ of a database schema is a set of table instances, one for each defined table (*extensional* relation) in the database. The notation $\Delta(T)$ represents the instance of a table T in Δ. Each query or view (*intensional* relation) R is defined as a multiset of facts, and Φ_R represents the ERA (Extended Relational Algebra) expression associated to an SQL query or view R, as explained in [11]. An intensional relation usually depends on previously defined relations, and sometimes it will be useful to write $\Phi_R(R_1, \ldots, R_n)$ indicating that R depends on $R_1, \ldots, R_n$. Here, we assume that each extensional relation in a database instance has attached type information for each one of its named arguments. As well, each intensional relation argument receives its type via inferencing and arbitrary names if not provided in its definition. Tables are denoted by their names, that is, $\Phi_T = T$ if T is a table.

Definition 1. The *computed answer* of an ERA expression Φ_R with respect to some schema instance Δ is denoted by $\parallel \Phi_R \parallel_\Delta$, where:

- If R is an extensional relation, $\parallel \Phi_R \parallel_\Delta = \Delta(R)$.
- If R is an intensional relation and $R_1, \ldots, R_n$ the relations defined in R, then $\parallel \Phi_R \parallel_\Delta = \Phi_R(\parallel \Phi_{R_1} \parallel_\Delta, \ldots, \parallel \Phi_{R_n} \parallel_\Delta)$. □

Queries are executed by SQL systems. The answer for a query Q and a database instance Δ in an implementation is represented by $\mathcal{SQL}_\Delta(Q)$. The notation $\mathcal{SQL}_\Delta(R)$ abbreviates $\mathcal{SQL}_\Delta(\texttt{SELECT * FROM } R)$. In particular, we assume the existence of *correct* SQL implementations.

Definition 2. A *correct* SQL implementation verifies that $\mathcal{SQL}_\Delta(Q) = \parallel \Phi_Q \parallel_\Delta$ for every query Q. □

3 Hypothetical Datalog

Hypothetical Datalog is an extension of function-free Horn logic [5]. Following [17], the syntax of the logic is first order and includes a universe of constant symbols, a set of variables and a set of predicate symbols. For concrete symbols, we write variables starting with either an upper-case letter or an underscore, and the rest of symbols starting with lower-case. Removing function symbols from the logic is a condition for finiteness of answers, a natural requirement of relational database users. As in Horn-logic, a rule has the form $A \leftarrow \phi$, where A is an atom and ϕ is a conjunction of goals. Since we consider a hypothetical system, a goal can also take the form $R_1 \wedge \ldots \wedge R_n \Rightarrow G$, a construction known as an *embedded implication*. The following definition captures the syntax of the language, where $vars(T)$ is the set of variables occurring in T:

Definition 3.
$R := A \mid A \leftarrow G_1 \wedge \ldots \wedge G_n$
$G := A \mid \neg G \mid R_1 \wedge \ldots \wedge R_n \Rightarrow G$
where R and R_i stand for rules, G and G_i for goals, A for an atom (possibly

containing variables and constants, but no compound terms), and $\bigcup vars(R_i) \cap vars(R) = \emptyset$, and $vars(R_i)$ and $vars(G)$ are disjoint. □

Disjoint conditions ensure that assumed rules do not depend on actual substitutions along inference, i.e., assumed rules take the form they have in the program.

Semantics is built with a stratified inference system which can be consulted in [17]. Here we recall the inference rule for the embedded implication:[2] For any goal ϕ and database instance Δ:

$$\frac{\Delta \cup \{R_1, \ldots, R_n\} \vdash \phi}{\Delta \vdash R_1 \wedge \ldots \wedge R_n \Rightarrow \phi}$$

This means that for proving the conclusion ϕ, rules R_i, together with the current database instance Δ can be used in subsequent inference steps. The *unified stratified semantics* defined in [17] builds a set of axioms $\mathcal{E}$ that provides a means to assign a meaning to a goal as: $solve(\phi, \mathcal{E}) = \{\Delta \vdash id : \psi \in \mathcal{E}$ such that $\phi\theta = \psi\}$, where θ is a substitution and each axiom in $\mathcal{E}$ is mapped to the database Δ it was deduced for, and the inferred fact ψ is labelled with its data source (for supporting duplicates). We use $\Delta(\mathcal{E})$ to denote the multiset of facts ψ so that $\Delta \vdash id : \psi \in \mathcal{E}$ for any id. So, this inference rule captures what SQL `WITH` statements need if translated to Hypothetical Datalog, because each R_i can represent each temporary view definition, as will be shown in the next section.

4 Translating SQL into Datalog

We consider standard SQL as found in many textbooks (e.g., [21]), but also allowing `FROM`-less statements, i.e., providing a single-row output constructed with the comma-separated expressions after the `SELECT` keyword (Oracle, for instance, resorts to feed the row from the `dual` table to express the same feature). Here, we define a function *`SQL_to_DL`* that takes a relation name and an SQL statement as input and returns a multiset of Datalog rules providing the same meaning as the SQL relation for a corresponding predicate with the same name as the relation. The following (incomplete) definition for this function includes only a couple of the basic cases, where others can be easily developed from [23]. From here on, set-related operators and symbols refer to multisets, as SQL relations can contain duplicates.

% Basic SELECT statement
`SQL_to_DL`$(r,$ `SELECT` $\mathtt{A}_1, \ldots, \mathtt{A}_n$ `FROM` Rel `WHERE` $Cond) =$
$\{\ r(\overline{X_i}) \leftarrow DLRel(\overline{X_i}), DLCond(\overline{X_j})\ \} \bigcup RelRules \bigcup CondRules,$
where *`SQLREL_to_DL`*$(Rel) = (DLRel(\overline{X_i}), RelRules)$, and
`SQLCOND_to_DL`$(Cond) = (DLCond(\overline{X_j}), CondRules)$

[2] Each rule in this inference system is read as: If the formulas above the line can be inferred, then those below the line can also be inferred.

% Duplicate-preserving union
$\mathit{SQL_to_DL}(r,\ \mathit{SQL}_1\ \texttt{UNION ALL}\ \mathit{SQL}_2) =$
$\quad \mathit{SQL_to_DL}(r,\ \mathit{SQL}_1) \bigcup \mathit{SQL_to_DL}(r,\ \mathit{SQL}_2)$

Here, each $\texttt{A}_i$ is an argument name present in the relation *Rel* with corresponding logic variable X_i. *Rel* is constructed with either a single defined relation (table or view), or a join of relations, or an SQL statement. Function $\mathit{SQLREL_to_DL}$ (resp., $\mathit{SQLCOND_to_DL}$) takes an SQL relation (resp. condition) and returns a goal and, possibly, additional rules which result from the translation. Variables $\overline{X_j}$ come as a result of the translation of the condition *DLCond* to a goal. As well, some basic cases are presented next for these functions, where $GoalName$ is an arbitrary, fresh new goal name:

% Extensional/Intensional Relation Name
$\mathit{SQLREL_to_DL}(RelName) = (RelName(\overline{X_i}), \{\})$
where $\overline{X_i}$ are the n variables corresponding to the n-degree relation $RelName$.

% SQL Statement
$\mathit{SQLREL_to_DL}(SQL) = (GoalName(\overline{X_i}),\ \mathit{SQL_to_DL}(GoalName, SQL))$
where $\overline{X_i}$ are the n variables corresponding to the n-degree statement SQL.

% NOT IN Condition
$\mathit{SQLCOND_to_DL}(\texttt{A}_i\ \texttt{NOT IN}\ Rel) = (not\ DLRel(\overline{X_j}), RelRules)$
where $\mathit{SQLREL_to_DL}(Rel) = (DLRel(\overline{X_j}), RelRules)$, and $X_i \in \overline{X_j}$ is the corresponding variable to argument $\texttt{A}_i$.

Completing this function by including the `WITH` statement is straightforward because every temporary view can be represented by a predicate resulting from the translation of the temporary view definition into Datalog rules. Assuming such predicates as the antecedent of an embedded implication can be used to augment the (local, temporary) database for interpreting the meaning of the translated SQL outcome query:

$\mathit{SQL_to_DL}(r, \texttt{WITH}\ r_1\ \texttt{AS}\ \mathit{SQL}_1\texttt{,}\ \ldots\texttt{,}\ r_n\ \texttt{AS}\ \mathit{SQL}_n\ \mathit{SQL}) =$
$\quad \{\ r(\overline{X_i}) \leftarrow$
$\quad\quad \wedge(\mathit{SQL_to_DL}(r_1,\mathit{SQL}_1)) \wedge \ldots$
$\quad\quad \wedge(\mathit{SQL_to_DL}(r_n,\mathit{SQL}_n)) \Rightarrow s(\overline{X_i})\ \} \bigcup \mathit{SQL_to_DL}(s,\mathit{SQL})$

where $\wedge(Bag)$ denotes $B_1 \wedge \cdots \wedge B_m$ $(B_i \in Bag)$.

The following theorem establishes the semantic equivalence of an SQL relation and its counterpart Datalog translation.

Theorem 1. The semantics of an SQL n-degree relation r defined by the query Q on a database instance Δ coincides with the meaning of a goal $r(\overline{X_i})$, $1 \leq i \leq n$, for $\Delta' = \Delta \bigcup \mathit{SQL_to_DL}(r,Q)$, that is: $\mathcal{SQL}_\Delta(Q) = \Delta(solve(r(\overline{X_i}), \mathcal{E}))$, where $\mathcal{E}$ is the unified stratified semantics for Δ'. □

5 Beyond the WITH Clause: Expressing Assumptions

As a novel feature, hypothetical SQL queries (absent in the standard) were introduced (inspired in [16]) in DES version 2.6 for solving "what-if" scenarios. Syntax for such queries is:

`ASSUME` SQL_1 `IN` Rel_1`, ...,` SQL_n `IN` Rel_n SQL`;`

which makes to *assume* the result of SQL_i in Rel_i when processing SQL. This means that the semantics of each Rel_i is either overloaded (if the relation already exists) or otherwise defined with the facts of SQL_i. Implementing this resorted to *globally* define each Rel_i, which is not the expected behaviour as its definition must be local to SQL. Roughly, solving an `ASSUME` query resorted to overload the meanings of each Rel_i (by inserting the required facts) before computing SQL and, after solving, to restore them (by deleting the same facts). This also precluded nested assumptions, and such statements were allowed only as top-level queries but not as part of query definitions. For instance, if it would be allowed, the following query would be incorrectly computed in that scenario:

```
ASSUME SELECT 1 IN r(a),
       (ASSUME SELECT 2 IN r(a) SELECT * FROM r) IN s
SELECT * FROM r,s;
```

because the meaning of `r` in the context of `SELECT * FROM r,s` would be overloaded with both `{(1)}` and `{(2)}`, instead of just with `{(1)}`.

Applying hypothetical reasoning in this case solves this issue, allowing us not only to use nested assumptions in both top-level queries and views, but also to take advantage of negative assumptions. A negative assumption allows to *remove* facts from the meaning of a relation, which broadens the applicability of queries in decision-support scenarios. To specify negative assumptions, `NOT IN` is used instead of just `IN`. Hypothetical Datalog in [19] introduces the notion of *restricted predicate* to handle negative assumptions in embedded implications. A restricted predicate includes at least a restricting rule whose head is an atom preceded by a minus sign. Its meaning is the set of facts deduced from regular rules minus the set of facts deduced from restricting rules. So, a negative assumption is modelled with a restricting rule in the antecedent of an embedded implication, so that the translation from SQL to Hypothetical Datalog for `ASSUME` statements becomes:

$$
\begin{aligned}
&\mathit{SQL_to_DL}(r, \texttt{ASSUME } \mathit{SQL}_1 \texttt{ [NOT] IN } r_1\texttt{,...,}\mathit{SQL}_n \texttt{ [NOT] IN } r_n \texttt{ } \mathit{SQL}) = \\
&\quad \{r(\overline{X_i}) \leftarrow \\
&\qquad \wedge(\mathit{SQL_to_DL}\,(r_1,\, \mathit{SQL}_1)[\texttt{[-]}r_1/r_1]) \wedge \ldots \\
&\qquad \wedge(\mathit{SQL_to_DL}\,(r_n,\, \mathit{SQL}_n)[\texttt{[-]}r_n/r_n]) \texttt{ => } s(\overline{X_i})\ \} \\
&\qquad \cup\ \mathit{SQL_to_DL}(s, \mathit{SQL})
\end{aligned}
$$

where $A[B/C]$ represents the application of the syntactic substitution C by B in all the rule heads in A, `[`T`]` represents that T is optional but it must occur if the corresponding i-th entry also occurs (i.e., if `NOT` occurs in the assumption for r_i, then `-` also occurs in the corresponding substitution).

6 Playing with the System

Translating an SQL query to Datalog in a practical system involves more features that the ones briefly suggested before and are out of the scope of this extended abstract. For example, the `SELECT` list can include expressions and scalar SQL statements, nested statements can be correlated, aggregate functions and grouping can be used, and so on. Also, from a capacity point-of-view, a needed stage in the translation is folding/unfolding of rules [22] to simplify the Datalog program resulting from the translation. This is quite relevant because deducing the meaning (either complete or restricted to a given call) of the involved relations along query solving is needed, therefore significantly augmenting their space and time requirements. Next, we introduce a couple of examples of this translation with the system DES [20], which in particular supports such features and inputs from several query languages, including Datalog and SQL. Here, we resort to the actual textual syntax of Datalog rules in this system, which follows the syntax of Prolog.

student
(adam)
(bob)
(pete)
(scott)

take
(adam,db)
(pete,db)
(pete,lp)
(scott,lp)

Let us consider a database containing the relations `student(name)` and `take(name, title)`. The first one states names of students and the second one the course (`title`) each student (`name`) is enrolled in. Types can be specified either with a Datalog assertion (as `:-type(student(name:string))` for the first case) or a DDL SQL statement (as `create table take(name string, title string)` for the second one, where a foreign key `take.name` → `student.name` could be stated as well). We consider the database instance depicted in the tables above.

The next SQL statement (looking for students that have not been already enrolled in a course) is translated as follows in a system session with DES 4.0:

```
DES> select * from student where name not in
        (select name from take)
Info: SQL statement compiled to:
  answer(A) :- student(A), not take(A,_B).
answer(student.name:string) ->
{ answer(bob) }
Info: 1 tuple computed.
```

This example shows a few of things. First, as a query Q is allowed at the system prompt, the call to the translation function becomes *SQL_to_DL*(`answer`,Q), i.e., the outcome relation is automatically renamed to the reserved keyword `answer`. Second, the outcome schema `answer(student.name:string)` shows that the single output argument comes from the argument `name` of the relation `student`, with type `string`. Third, following the definition of the translation function, this query should be translated into:

```
answer(A) :- student(A), goal1(A).
goal1(A) :- not take(A,_B)
```

But folding/unfolding simplifies this as it was displayed in the system session. These translations have been displayed because it was specified so by issuing the command `/show_compilations on`. Finally, non-relevant variables to a rule outcome are underscored (otherwise, they would be signalled as anonymous). This is important in this case to identify as safe the rule in which this underscored variable occurs. Classical safety [23] would tag the rule `answer(A) :- student(A), not take(A,_B)` as unsafe, but an equivalent set of safe rules can be found: `answer(A) :- student(A), not goal1(A)` and `goal1(A) :- take(A,_B)`. Underscored variables are a means to encapsulate this form of safety, which is identified by the system and processed correspondingly. In general, there can be several rules for `answer` (e.g., when a `UNION` is involved) and others on which this predicate depends on.

As an example of a `WITH` query, the following statement defines the relation `grad` intended to retrieve the eligible students for graduation (those that took both `db` and `lp` in this tiny example):

```
DES> with grad(name) as
       (select student.name
        from student, take t1, take t2
        where student.name=t1.name
          and t1.name=t2.name
          and t1.title='db' and t2.title='lp')
     select * from grad;
Info: SQL statement compiled to:
  answer(A) :-
  (grad(B) :- student(B), take(B,db), take(B,lp))
  =>
  grad(A).
answer(grad.name:varchar(30)) ->
{ answer(pete) }
Info: 1 tuple computed.
```

As an example of an `ASSUME` query, we reuse the `grad` definition above, assume that `adam` is not an eligible student, and that `adam` and `scott` took `lp` and `db` respectively:

```
DES> assume
       (select 'adam') not in student,
       (select 'adam','lp' union all select 'scott','db')
         in take,
       (select student.name from student, take t1, take t2
         where student.name=t1.name and t1.name=t2.name and
               t1.title='lp' and t2.title='db') in grad(name)
     select * from grad;
Info: SQL statement compiled to:
  answer(A) :-
  -student(adam) /\ take(adam,lp) /\  take(scott,db) /\
  (grad(B) :- student(B), take(B,lp), take(B,db))
  =>
```

```
  grad(A).
answer(grad.name:varchar(30)) ->
{ answer(pete), answer(scott) }
Info: 2 tuples computed.
```

Here, the assumption on `student` is negative and is compiled to a restricting fact. The second one is compiled to a couple of facts because of the union. The last one is the same as the previous example. The SQL statement after the assumptions simply leads to the goal `grad(A)`, for which even when `adam` took the courses to graduate, he was removed as an eligible student and therefore from the answer.

If the extensional relations `student` and `take` are already defined in an external relational database (as, e.g., MySQL or PostgreSQL), they can be made available to DES via an ODBC connection (with the command `/open_db`), and queried as if they were local [18]. This way, DES behaves as a front-end for both straight calls to native (i.e., supported by the external relational system) SQL queries and non-native queries (as those including `ASSUME`). For non-native statements, prepending the command `/des` to the query makes DES to handle such queries which are unsupported in the external database. For example:

```
DES> /open_db postgresql
DES> /des assume ...
DES> /open_db mysql
DES> /des with ...
```

obtaining the same answers as before for the same queries (both omitted here in the ellipses). Note that, in particular, `WITH` is unsupported in both MySQL and MS Access.

Even when `WITH` is supported in several relational database systems, they are somewhat restricted because, referring to the syntax in Section 2, *SQL* cannot contain a `WITH` clause, whereas we do allow for it.

7 Conclusions

This work has presented a proposal to take advantage of intuitionistic logic programming to model both temporary definitions (with the `WITH` clause) and assumptions (with the `ASSUME` clause) in SQL. Its motivation lies in providing a clean semantics that makes assumptions to behave as first-class citizen in the object language. The deductive database system DES was used as a test bed to experiment with assumptions, translating SQL queries into Hypothetical Datalog. Further, this system can be used as a front-end to relational systems lacking features as the `WITH` clause. The most related work is [1], which includes assumptions in SQL with a tailored semantics, and generates SQL scripts implementing fixpoint computations. With respect to the intuitionistic formal framework, our work is based on [3,4,5,6] and adapted to assume rules and deal with duplicates in [17]. However, it is not powerful enough to include embedded universal quantifiers in premises as in [6], which provides the ability to create new constant

symbols hypothetically along inference. Though this is not directly applicable to the current work, it is indeed an interesting subject to explore by considering that domains can be finitely constrained in practical applications, as with foreign keys.

Acknowledgements

Thanks to the anonymous referees for their suggestions to improve this work, which has been partially supported by the Spanish MINECO project CAVI-ART (TIN2013-44742-C4-3-R), Madrid regional project N-GREENS Software-CM (S2013/ICE-2731) and UCM grant GR3/14-910502.

References

1. G. Aranda-López, S. Nieva, F. Sáenz-Pérez, and J. Sánchez-Hernández. Incorporating Hypothetical Views and Extended Recursion into SQL Database Systems. In K. Mcmillan, A. Middeldorp, G. Sutcliffe, and A. Voronkov, editors, *International Conference on Logic for Programming, Artificial Intelligence and Reasoning (LPAR-19)*, volume 26 of *EPiC Series in Computing*, pages 9–22. EasyChair, 2014.
2. P. Atzeni, C. S. Jensen, G. Orsi, S. Ram, L. Tanca, and R. Torlone. The Relational Model is Dead, SQL is Dead, and I Don't Feel So Good Myself. *SIGMOD Record*, 42(2):64–68, July 2013.
3. A. J. Bonner. Hypothetical Datalog: Negation and Linear Recursion. In *Proceedings of the ACM Symposium on Principles of Database Systems (PODS)*, pages 286–300, 1989.
4. A. J. Bonner. Hypothetical Datalog: Complexity and Expressibility. *Theoretical Computer Science*, 76:3–51, 1990.
5. A. J. Bonner and L. T. McCarty. Adding Negation-as-Failure to Intuitionistic Logic Programming. In E. L. Lusk and R. A. Overbeek, editors, *Proceedings of the North American Conference on Logic Programming (NACLP)*, pages 681–703. The MIT Press, 1990.
6. A. J. Bonner, L. T. McCarty, and K. Vadaparty. Expressing Database Queries with Intuitionistic Logic. In E. L. Lusk and R. A. Overbeek, editors, *Proceedings of the North American Conference on Logic Programming (NACLP)*, pages 831–850, 1989.
7. R. Caballero, Y. García-Ruiz, and F. Sáenz-Pérez. Declarative Debugging of Wrong and Missing Answers for SQL Views. In *Eleventh International Symposium on Functional and Logic Programming (FLOPS)*, LNCS 7294. Springer, 2012.
8. E. Codd. A Relational Model for Large Shared Databanks. *Communications of the ACM*, 13(6):377–390, June 1970.
9. E. Codd. Relational Completeness of Data Base Sublanguages. In Rustin, editor, *Data base Systems*, Courant Computer Science Symposia Series 6. Englewood Cliffs, N.J. Prentice-Hall, 1972.
10. D. M. Gabbay. N-Prolog: An Extension of Prolog with Hypothetical Implication II - Logical Foundations, and Negation as Failure. *Journal of Logic Programming*, 2(4):251–283, 1985.
11. H. Garcia-Molina, J. D. Ullman, and J. Widom. *Database Systems: The Complete Book*. Prentice Hall PTR, Upper Saddle River, NJ, USA, 2008.

12. P. W. Grefen and R. A. de By. A Multi-Set Extended Relational Algebra: A Formal Approach to a Practical Issue. In *Proceedings of the Tenth International Conference on Data Engineering (ICDE)*, pages 80–88. IEEE, 1994.
13. J. Hodas and D. Miller. Logic Programming in a Fragment of Intuitionistic Linear Logic. *Information and Computation*, 110(2):327 – 365, 1994.
14. L. T. McCarty. Clausal Intuitionistic Logic I - Fixed-Point Semantics. *Journal of Logic Programming*, 5(1):1–31, 1988.
15. D. Miller. A Logical Analysis of Modules in Logic Programming. *The Journal of Logic Programming*, 6(1):79 – 108, 1989.
16. S. Nieva, F. Sáenz-Pérez, and J. Sánchez. Formalizing a Constraint Deductive Database Language based on Hereditary Harrop Formulas with Negation. In *FLOPS'08, Proceedings*, volume 4989 of *LNCS*, pages 289–304, Ise, Japan, 2008. Springer-Verlag.
17. F. Sáenz-Pérez. Implementing Tabled Hypothetical Datalog. In *Proceedings of the 25th IEEE International Conference on Tools with Artificial Intelligence (ICTAI)*, pages 596–601, November 2013.
18. F. Sáenz-Pérez. Towards Bridging the Expressiveness Gap Between Relational and Deductive Databases. *Electronic Communications of the EASST*, 64:1–22, 2014.
19. F. Sáenz-Pérez. Restricted Predicates for Hypothetical Datalog. *Electronic Proceedings in Theoretical Computer Science*, 200(0):64–79, 2015.
20. F. Sáenz-Pérez. Datalog Educational System, 2016. `http://des.sourceforge.net/`.
21. A. Silberschatz, H. Korth, and S. Sudarshan. *Database Systems Concepts*. McGraw-Hill, Inc., New York, NY, USA, 6th edition, 2010.
22. L. Sterling and E. Shapiro. *The Art of Prolog: Advanced Programming Techniques*. MIT Press, Cambridge, MA, USA, 1986.
23. J. D. Ullman. *Database and Knowledge-Base Systems, Vols. I (Classical Database Systems) and II (The New Technologies)*. Computer Science Press, 1988.